Abel inversion using total-variation regularization

Thomas J. Asaki, asaki@lanl.gov
Patrick R. Campbell, prc@t7.lanl.gov
Rick Chartrand, rickc@lanl.gov
Collin Powell, cep@lanl.gov
Kevin R. Vixie, vixie@lanl.gov
Brendt Wohlberg, brendt@lanl.gov

The density reconstruction of a cylindrically symmetric object from a single radiographic view is a classic and important tomography problem. It comes up in astronomy, the study of plasmas and flames, and material property studies. In principle, it has a straightforward solution: if a beam of radiation, in which all the rays are parallel, passes through such an object perpedicularly to the axis of symmetry, then the attenuation of the radiation is given by the Abel transform. Thus, the radiograph has a simple, mathematical description in terms of the density at each point of the object. Moreover, this description is invertible: the object's density is given by the inverse Abel transform of the attenuation measured by the radiograph. In principle, this allows one to calculate from the radiograph the density at each point of the object.

Unfortunately, this approach is of limited usefulness. The reason is that the inverse Abel transform will magnify noise present in the radiograph, giving a very noisy reconstruction.

One way around this is to *regularize* the reconstruction. Let d be the function giving the density of the object at each point, r the attenuation measured by the radiograph, A the Abel transform. The goal is to determine d from r, with $d = A^{-1}r$ being unsatisfactory. Instead, we choose a regularization functional R and a data-fidelity functional F, and seek to minimize the quantity

$$R(d) + F(Ad - r).$$

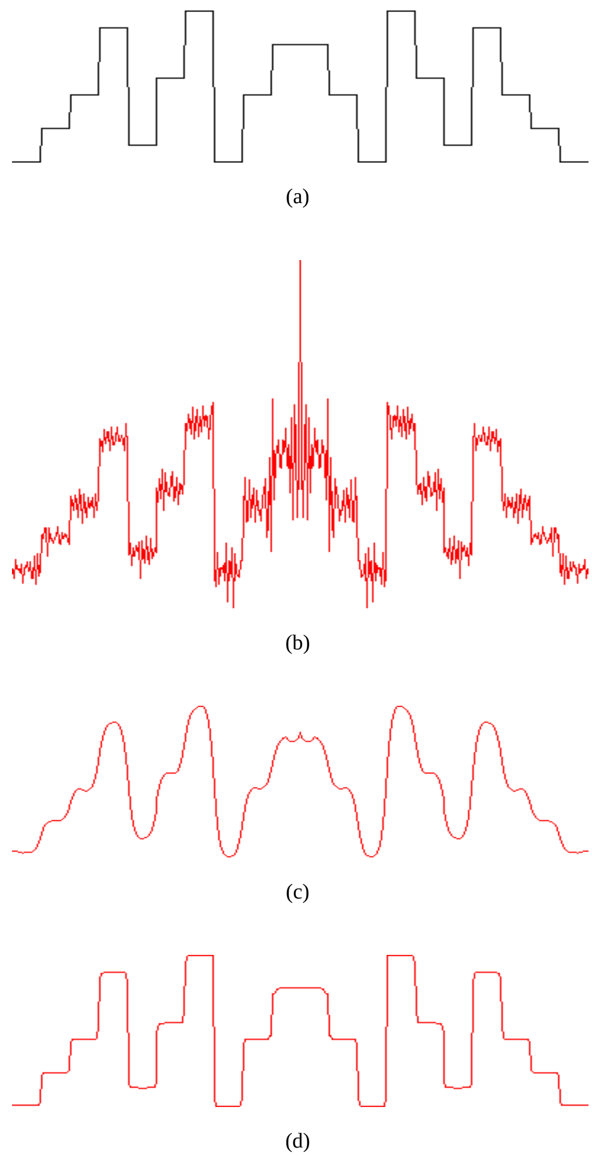

(a) (b) (c) (d)

Comparing Abel inversions using different regularizations. (a) Density profile of the cross section of a two-dimensional, cylindrically symmetric object. A simulated, noisy radiograph of this object is produced, and different Abel inversion methods applied. (b) Unregularized Abel inversion. Radiograph noise is amplified by the inverse Abel transform. (c) H^1 regularization. This regularization forces a smoothness that causes edge information to be lost. (d) TV regularization. Noise is suppressed, but edges are preserved.

A common choice for F is the square of the L^2 norm, so that

$$F(Ad-r) = \int |Ad-r|^2.$$

A traditional choice for R is the H^1 seminorm:

$$R(d) = \alpha \int |\nabla d|^2,$$

where ∇ is the gradient operator and α is a parameter to be chosen. By finding d that minimizes the sum of these two terms, d will be a smooth function (to keep $R(d)$ small) whose Abel transform is close to r (to keep $F(Ad-r)$ small). The size of the parameter α dictates the relative importance of the two terms, and should be chosen so that the data fidelity term will end up being equal to the estimated variance of the noise. The result is an approximate Abel inverse having very little noise.

Unfortunately, this method is unsuitable when the object has sharp changes in density, such as at the boundary between two different materials. The smoothness of d means that its value cannot change sharply, and information about material boundaries will be lost. To remedy this, we replace the regulzarization term with the total variation seminorm,

$$R(d) = \alpha \int |\nabla d|.$$

The simple act of replacing the exponent 2 with 1 allows discontinuous functions d to be solutions. The regularization term still suppresses noise, but while allowing sharp material boundaries. The improved performance of this method can be seen in the figure.

This approach comes with the cost of being much more difficult to implement computationally. A simple algorithm for computing the d that minimizes $R(d)+F(Ad-r)$ is the method of gradient descent. However, gradient descent requires many iterations to converge to a solution. Instead, we implemented a form of quasi-Newton method designed particularly for total-variation regularized inverse problems, the lagged diffusivity method of Vogel and Oman. Another improvement was to use the unregularized Abel inverse as the initial value of d with which to begin the iteration. The result is an algorithm that converges in only a few iterations.

Two papers on this work have been submitted for publication. One concerns the theoretical underpinnings of the optimization algorithm, the other the numerical implementation and the application to radiographic data.

Acknowledgements

Los Alamos Report LA-UR-05-0680. Funded by the Department of Energy under contracts W-7405-ENG-36 and the DOE Office of Science's Advanced Scientific Computing Research (ASCR) program in Applied Mathematical Sciences. The authors acknowledge the support from the Advanced Simulation and Computing (ASC) Program.